

Modernizing Fortran 77 Legacy Codes

Charles D. Norton¹ and Viktor K. Decyk^{1,2}

High Performance Computing Systems and Applications Group and
Center for Space Mission Information and Software Systems

¹ National Aeronautics and Space Administration
Jet Propulsion Laboratory, California Institute of Technology
MS 168-522
4800 Oak Grove Drive
Pasadena, California 91109-8099

² Department of Physics and Astronomy
University of California at Los Angeles
Los Angeles, CA 90095-1547

Email: Charles.D.Norton@jpl.nasa.gov, Viktor.Decyk@physics.ucla.edu

Keywords: Fortran 90, Scientific Programming, Legacy Codes

1 Modernizing Scientific Software

Legacy software has great value since it is generally well debugged, produces results that are trusted, and is actively meeting end-user goals. The amount of hidden expert knowledge embedded in such software can be significant making its preservation important. Nevertheless, legacy software has limitations. It can be difficult to extend, modify, and it does not support collaborative development very well. This can impede the ability to meet new and expanded mission goals as timelines and budgets become tighter. One approach to this problem is to rewrite the software from scratch, but this may introduce more serious costs. In particular, developing new verification and validation tests can be expensive. Also, ensuring that the legacy code was faithfully rewritten, regardless of the programming language applied, cannot always be guaranteed.

Generally, if the functionality of the legacy software is sound, it can be wrapped in a modern interface where the original code is mostly unmodified. The idea of wrapping code means that the original legacy software is preserved while a new layer of software is introduced to separate the old software from the new software. The wrapper provides the best means of retaining the functionality of the legacy software investment while providing a more flexible context from which new software, based on modern concepts, can be introduced.

There are many benefits to this approach:

1. Software remains in productive use while applications are modernized.

2. Avoids costly and potentially harmful software rewrites.
3. Promotes collaborative development while resolving organization problems exhibited in older codes.
4. Re-engineering occurs more quickly than rewriting, while preserving verification and validation tests, especially when the original programmers are involved.
5. Old bugs are uncovered.


Extending the functionality of legacy systems has become more important as modern applications increase in complexity and require the interaction of multiple contributors.

1.1 The Technology Applied

We have found that Fortran 90/95 has new features to support object-oriented principles beneficial for scientific programming and introduce a design methodology that defines a step-by-step process to modernize legacy application codes using state-of-the art software practices. While we emphasize Fortran applications, due to the abundance of Fortran legacy codes, similar techniques can be applied to software written in other domain-specific or general-purpose languages, including C or C++. As such, existing flight-related software is not excluded from modernization.

Our process begins by upgrading the existing Fortran application to standard conforming Fortran 90/95. Next, interfaces to the original application routines are introduced to add safety features by detecting common programming errors. These interfaces ensure that the

wrapper layer, added next, always correctly calls the legacy code. This wrapper layer allows problem based object abstractions to be introduced that interact cleanly with the legacy code, while supporting new enhancements. It also preserves the original, mostly unmodified, legacy software. The user can communicate with the modernized code across these layers and continuous development can occur simultaneously among these layers. The figure below shows the general approach



1.2 Fortran 90/95 Features that Modernize Programming

Many new features in the Fortran 90/95 standard provide benefits that are unfamiliar to experienced Fortran 77 software engineers. These features add safety, simplify complex operations, and allow software to be organized in a logically related way. Since backward compatibility is preserved one can incrementally make modifications while preserving existing work. Briefly, some of these new features are:

Modules Encapsulates (groups together) data, routines, and type declarations while providing accessibility across program units.

Use-Association Controls access to module content across program units.

Interfaces Verifies that the argument types in the procedure call match the types in the procedure declaration.

Derived Types User-defined types that support abstractions in programming. The creation of these types allows one to support problem domain based design.

Array Syntax This syntax simplifies whole array, and array subset, operations.

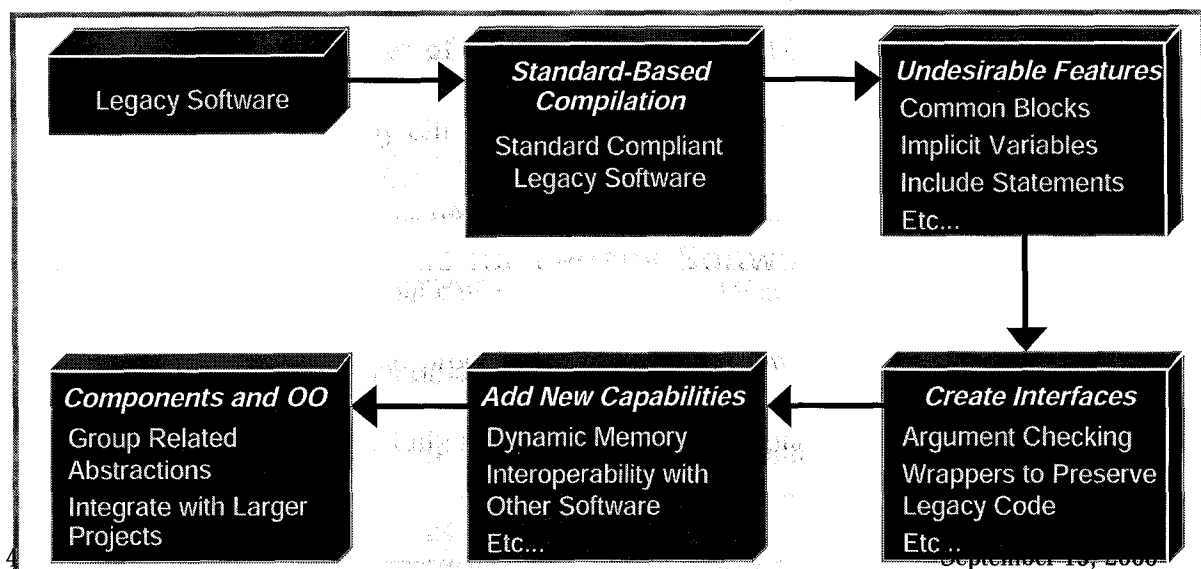
Dynamics Various kinds of dynamic structures are supported including allocatable arrays and pointers.

A very powerful realization is that combining these ideas allows support for object-oriented concepts. There are a number of textbooks on Fortran 90. One that we recommend is "Fortran 90 Programming", by Ellis, Philips, and Lahey, Addison Wesley, 1994.

2 Step-by-Step Process for Legacy Software Modernization

In this section we describe the process that one can follow to modernize legacy software.

This process has been successfully applied and it defines a plan of action for such projects.



The diagram shows the fundamental stages involved. While we focus on Fortran legacy codes the same stages could be modified for software written in other languages. Many of the specific actions taken will also depend on the code structure and objectives.

2.1 Clearly Identify the Objectives

It is very important to have a conversation with the software owners to determine their objectives. The flowchart of the modernization process may help guide this discussion.

2.2 Understand the Legacy Software

Understanding, even at a basic level, how the legacy software is organized is valuable. While it is possible to perform the modernization without detailed knowledge of the application, knowing the design is very helpful. Here are some common questions that should be asked.

- Is this a stand-alone code or is additional software required?
- Is this a single language code or a multilanguage code?
- What platforms are required?
- Who is responsible for answering questions if legacy bugs are detected?
- What kind of obsolete features exist in the software?
- Are any third-party developers involved and is their software proprietary?

2.3 Addressing Undesirable Features

One of the most undesirable features in legacy Fortran 77 codes are COMMON blocks since they often inhibit more advanced features, like dynamic memory. They also discourage code sharing since everything is exposed. For this reason, modifying large common blocks can also

be intimidating since inadvertent errors are easy to introduce. Other undesirable features include implicitly declared variables, which are dangerous, and include statements that are platform dependent based on how directories are specified.

Common blocks can be handled by placing the specification in a Fortran 90/95 module. Furthermore, rather than using `include` to make a textual substitution, the module information can become accessible using the Fortran 90/95 `use` statement in the appropriate routines.

```
! Original COMMON Block in common.inc
real arg1(10,10), arg2(10,10)
logical arg3
integer arg4
COMMON /BLOCK1/ arg1,arg2,arg3,arg4
SAVE /BLOCK1/

subroutine foo()                ! Defined in some file...
include 'common.inc'
...
end

! Modernized Version in common.f
MODULE common_block1
  implicit none
  save
  real, dimension(10,10) :: arg1, arg2
  logical :: arg3
  integer :: arg4
END MODULE common_block1

subroutine foo()                ! Defined in some file...
  use common_block1
...
end subroutine foo
```

The structure of the replacement is straightforward. One could have simply copied the original common block from `common.inc` into a module exactly, but using the Fortran 90/95

constructs gives additional advantages. These include the ability to make the block members dynamic and the ability to add more functionality to the module by making other modules visible within its scope, to name a few.

2.4 Creating Interfaces

Interfaces are very important, as they add safety to the modernized software. They allow the compiler to verify consistent argument usage for procedures, which allows subtle errors to be detected and corrected in legacy codes.

Interfaces are created automatically for routines that are defined within modules, but we are currently interested in building interfaces for the legacy routines that will not be moved into modules at this time. Not every legacy routine requires an interface, but all of the routines accessible from the main program should have an interface. Furthermore, any routines in the scope of the main program that have arguments that will be dynamic will require an interface.

The `interface` statement is used to declare the procedure name and the types of its arguments. Since this is a Fortran 90/95 construct that will tie in the legacy code to the modernized code a new Fortran 90/95 `interface.f` file can be created to declare the Fortran 77 legacy interfaces. These interfaces can be placed in a module that in turn may use other modules, such as the common block modules recently created.


```

! Interface Module in interface.f
MODULE interface_module
  USE common_block1
  implicit none
  save
  interface
    subroutine foof77(arg1, arg2, dim1)
      real arg1(dim1,dim1)
      integer dim1
      logical arg2
    end subroutine
  end interface
END MODULE interface_module

```

Note that the interface has exactly the same declaration as the original Fortran 77 legacy procedure, in fact it is best to just copy it explicitly. This means that when the legacy routine is called additional checks will be performed to ensure that the number and types of the arguments match exactly.

It may look like very little has been gained, but the benefit of the interface becomes clear when it is combined with a wrapper that allows more powerful Fortran 90/95 features to be applied. For example, many Fortran 77 programs have very long argument lists because extra information must be included, such as the dimension of arrays. Since Fortran 90/95 arrays know their size these arguments do not need to be included in a wrapper function that calls the original legacy procedure.

```

! Interface Module in interface.f
MODULE interface_module
  USE common_block1
  implicit none
  save
  interface
    subroutine foof77(arg1, arg2, dim1)
      real arg1(dim1,dim1)
      integer dim1
      logical arg2
    end subroutine
  end interface
CONTAINS
  subroutine foof90(arg1, arg2)                                ! Wrapper
    real, dimension(:,:) :: arg1
    logical :: arg2
    call foof77(arg1, size(arg1,1), arg2)                      ! Legacy
  end subroutine foof90
END MODULE interface_module

```

This is a simple example, but the effect can be significant for very complex procedures. In fact, more functionality (such as dynamic memory) can be applied at this level using the wrapper while preserving the original legacy software. Furthermore, this can be achieved without a serious performance penalty when the legacy routine is non-trivial.

The interfaces can also clarify how Fortran 77 style arguments are sometimes passed to procedures. For example, it is not uncommon to find Fortran 77 programs that pass a two-dimensional array to a procedure that expects a one-dimensional array. This can cause compile errors when interfaces are used because they require that the arguments must match exactly. In such instances it is possible to create multiple interfaces to recognize this difference using a generic procedure to allow a single name to select the correct module procedure based on the argument list.

2.5 Adding New Capabilities

Now that the interfaces have been created and wrappers have been introduced to encapsulate the legacy software new capabilities can be added. For most legacy software the most desirable feature is dynamic memory. Fortran 90/95 supports many kinds of allocatable structures and they are straightforward to use. Dynamic memory increases the flexibility of the software since this frees the application user from fixed problem sizes. Interoperability with more modern software can also be achieved since the wrappers can be designed to utilize such applications. These new capabilities can be added without affecting the use of existing systems.

```
! Legacy Fortran 77 include file of static COMMON data
parameter (mdttl=128)
integer nElt, RayID(mdttl,mdttl), ...
COMMON /EltInt/ nElt, RayID,...
SAVE /EltInt/

! New Module for COMMON data
MODULE elt_common
  implicit none
  save
  integer :: nElt, mdttl = 128
  integer, allocatable, dimension(:, :) :: RayID
CONTAINS
! Constructor
  subroutine new_elt_common()
    allocate( RayID(mdttl,mdttl) )
  end subroutine new_elt_common
END MODULE elt_common

! Dynamic allocation from main program
PROGRAM example
  use elt_common
  implicit none
  call new_elt_common()
  ...
END PROGRAM example
```

The example above shows a legacy Fortran 77 common block with static data can be reorganized to support dynamic memory. This occurs by moving the common block into a module and specifying which structures will be dynamic. A constructor can be created to perform the allocation of the dynamic structure and this constructor can be called from the main program. A number of additional safety features such as checking if the structure was already allocated, handling of exceptional conditions like insufficient memory, and so forth can be added as well.

2.6 Moving toward Components and Object-Oriented Design

Fortran 90/95 contains derived types, like structures in C, which allow users to create their own types. This allows one to program using designs that better represent the problem domain. One of the major benefits of the methodology is that one can incrementally evolve the legacy code toward such a design while preserving the functionality of the legacy software. An object-oriented design allows the implementation details to change without affecting the user. In a sense, the interfaces and wrappers have hidden the details of the legacy software, but we can enhance the wrappers to support derived types evolving the code toward an object-oriented, component-based, design.

```

! Creating derived types for object-based design
type species
  real, dimension(:,:), pointer :: coords
  real :: charge_to_mass, kinetic_energy
end type species

! Using a legacy routine through an OO wrapper
subroutine w_push(particles, force, dt)
type (species) :: particle
real, dimension(:) :: force
real :: dt, qbm, wke
integer :: ndim, nparticle, nx
  ndim = size(particle%coordinates,2)
  nx = size(force)
  qbm = particle%charge_to_mass
  wke = particle%kinetic_energy
  call push(particle%coords, force, qbm, wke, ndim,
            nparticle, nx, dt)
end subroutine w_push

```

This example shows a legacy push(...) routine, wrapped by a Fortran 90/95 w_push(...), routine that uses a derived type to group together related information. This was not possible in Fortran 77 so long complicated argument lists were required. The species type has a dynamic component, and other information, which simplifies the programmer's view of the data. Nevertheless, the original legacy software can still provide the functionality required.

```

! Creating classes for object-based design
MODULE plasma_class
! Create Derived Types...
CONTAINS
  subroutine new_species(...) ! Constructor...
  subroutine w_push(...) ! Class Members...
END MODULE plasma_class

```

In fact, a class can be created which groups together operations common to the new species type where the class member routines utilize legacy software internally. New software can be

added to the class as well. This is a very powerful concept, but careful planning is always required when building an object-oriented design.

Once the classes have been designed, and tested, the modules can be incorporated into the main program and calls to the member routines can replace calls to the original legacy software. Since interfaces for the legacy software still exist this process can be incremental, the software still works at the end of the day, and development can continue during the modernization process allowing existing objectives to be satisfied.

3 Experiences and Comments

Legacy software still has great value, but extending that functionality has become more important. Modern applications require greater complexity and support for multiple authors. More flexible design and dynamic features are also beneficial. Our methodology allows a modern superstructure to be erected around a legacy code where data abstraction and information hiding help to limit exposure of unnecessary details. Furthermore, modern software features, combined with improvements in compilers, help to reduce inadvertent errors. Applying wrappers to protect legacy software allows one to extend the functionality of that software. The re-engineering methodology's cost effectiveness, speed, and ability to protect one's investment in the experience and knowledge embedded within legacy software should be beneficial to a variety of mission software projects.

Acknowledgment

This work was developed at the Jet Propulsion Laboratory, California Institute of Technology under a contract with the National Aeronautics and Space Administration and supported by a grant from the Jet Propulsion Laboratory Center for Space Mission Information and Software Systems Software Engineering Technology Work Area. The work of V. K. D. was also partially supported by US Department of Energy, Office of Fusion Science.